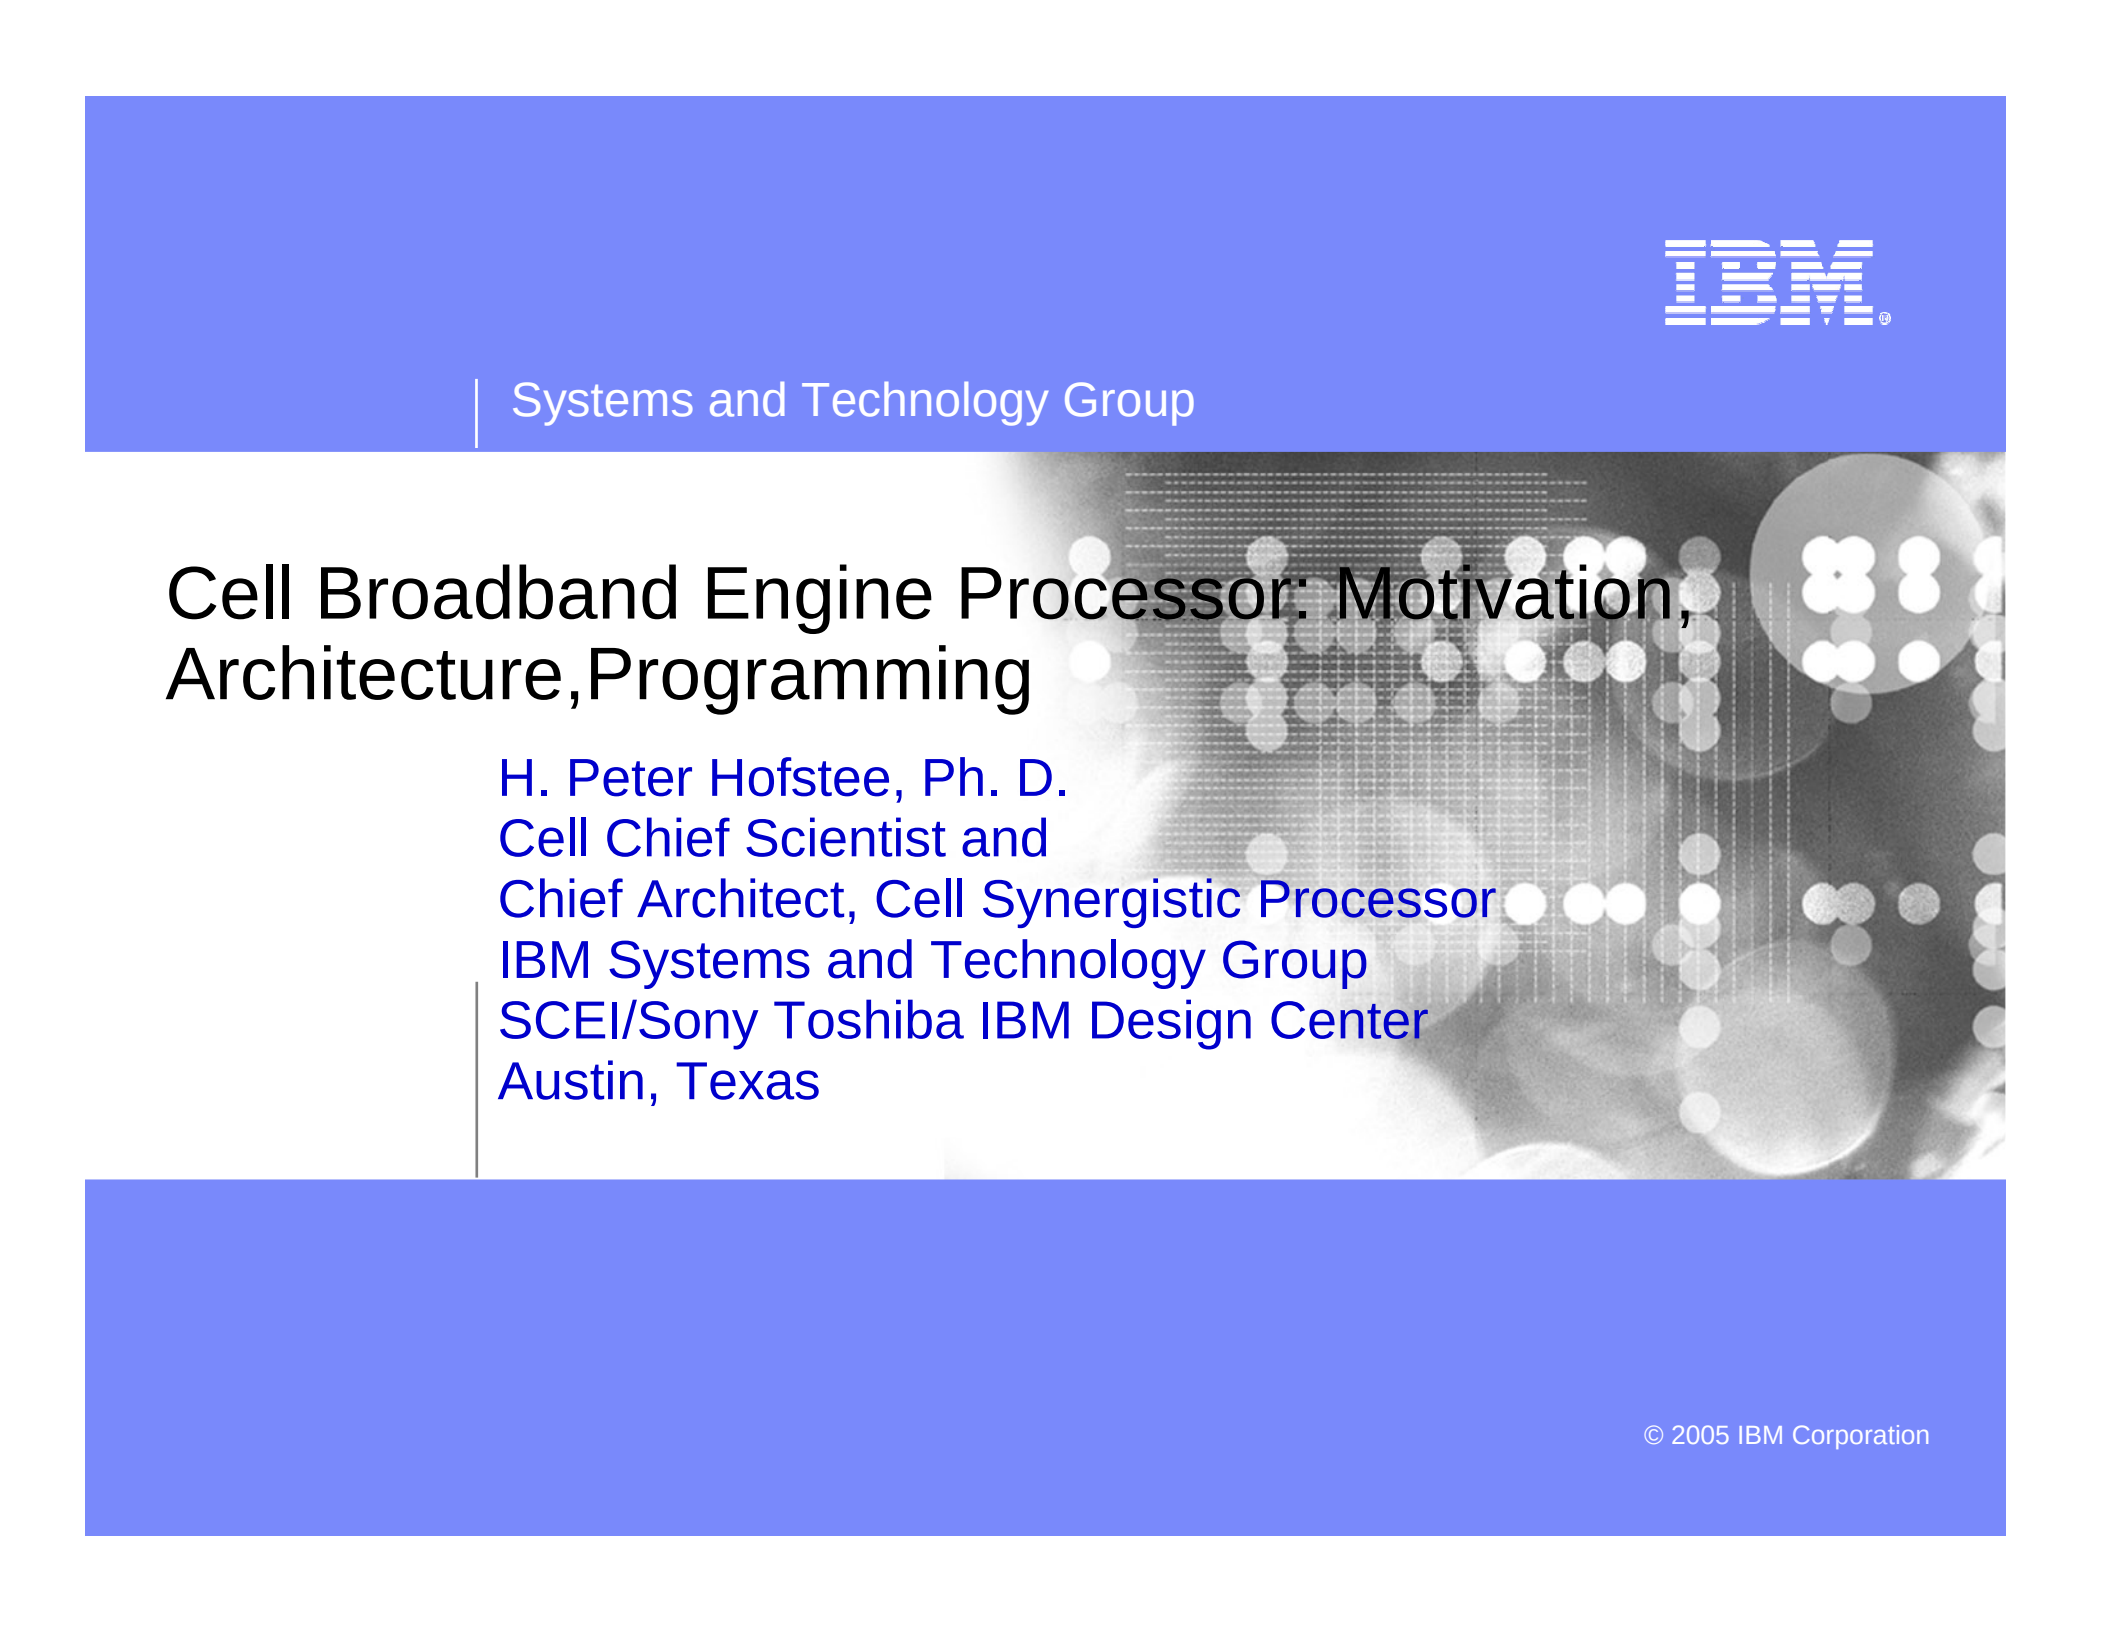

Cell Broadband Engine Processor: Motivation, Architecture,Programming

H. Peter Hofstee, Ph. D.
Cell Chief Scientist and
Chief Architect, Cell Synergistic Processor
IBM Systems and Technology Group
SCEI/Sony Toshiba IBM Design Center
Austin, Texas

Acknowledgements

- **Cell Broadband Engine (“Cell”) is the result of a deep partnership between SCEI/Sony, Toshiba, and IBM**
- **Cell represents the work of more than 400 people starting in 2001and a design investment of about $400M**

Agenda

- **Motivation**
- **Architecture**
- **Implementation**
- **Programming**
- **Applications**

Motivation

Motivation: Cell Goals

- **Outstanding performance, especially on game/multimedia applications.**
  - Challenges: Power Wall, Frequency Wall, Memory Wall
- **Real time responsiveness to the user and the network.**
  - Challenges: Real-time in an SMP environment, Security
- **Applicable to a wide range of platforms.**
  - Challenge: Maintain programmability while increasing performance
- **Support an introduction in 2005/6.**
  - Challenge: Structure innovation such that 5yr. schedule can be met

Power Wall: Module Heat Flux Trend

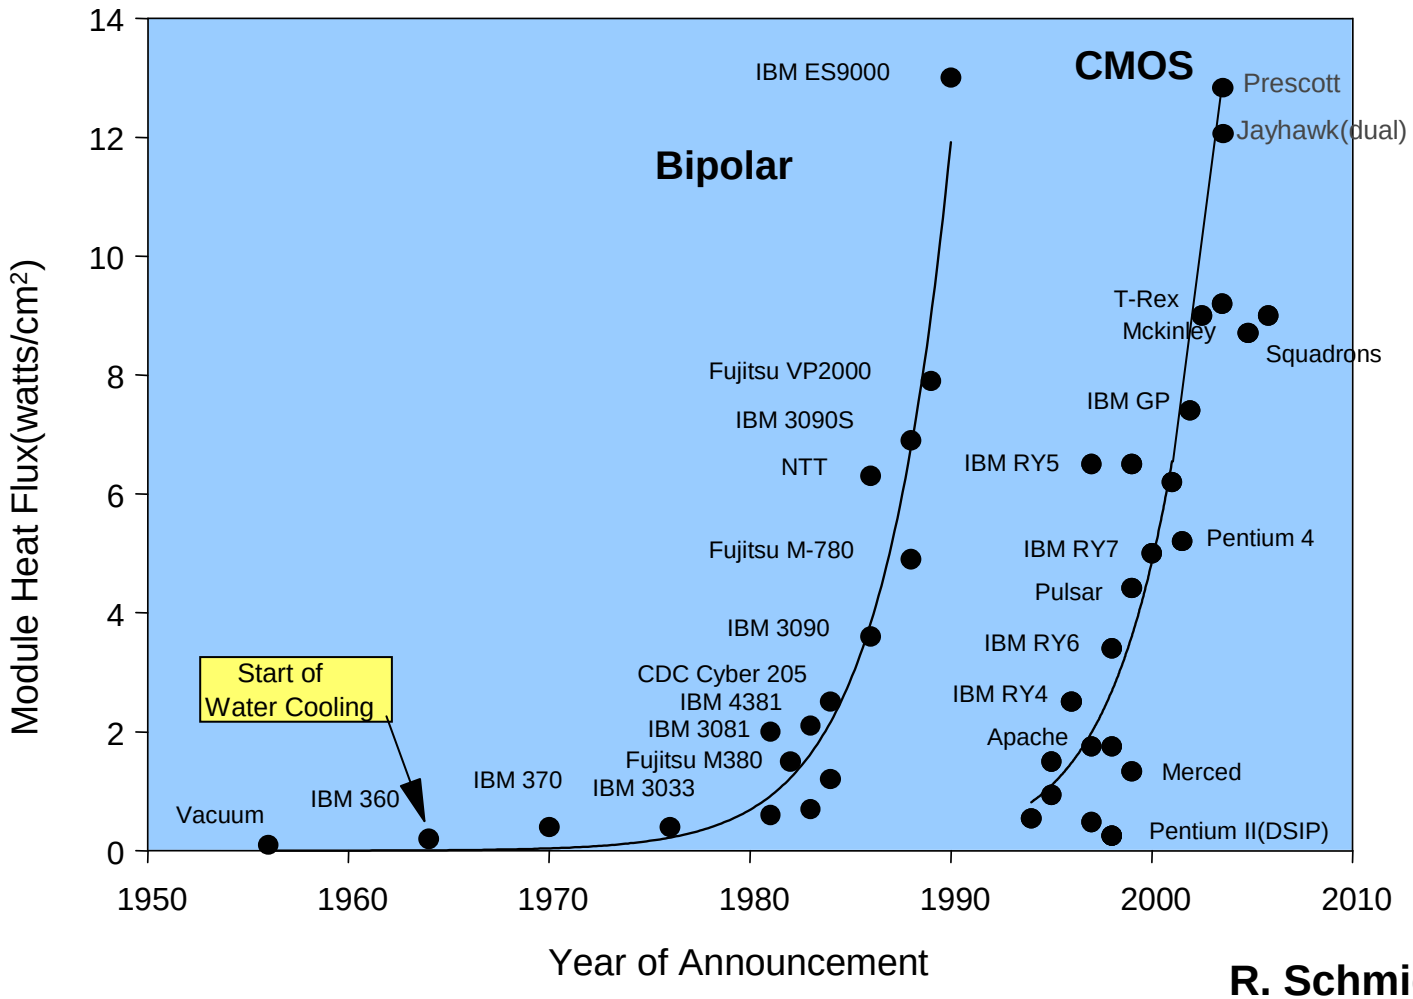

R. Schmidt, IBM

CMOS Devices hitting a scaling wall

- **Power components:**
  - Active power
  - Passive power
    - Gate leakage
    - Sub-threshold leakage (source-drain leakage)
- **Net:**
  - Further improvements require structure/materials changes (next slide)

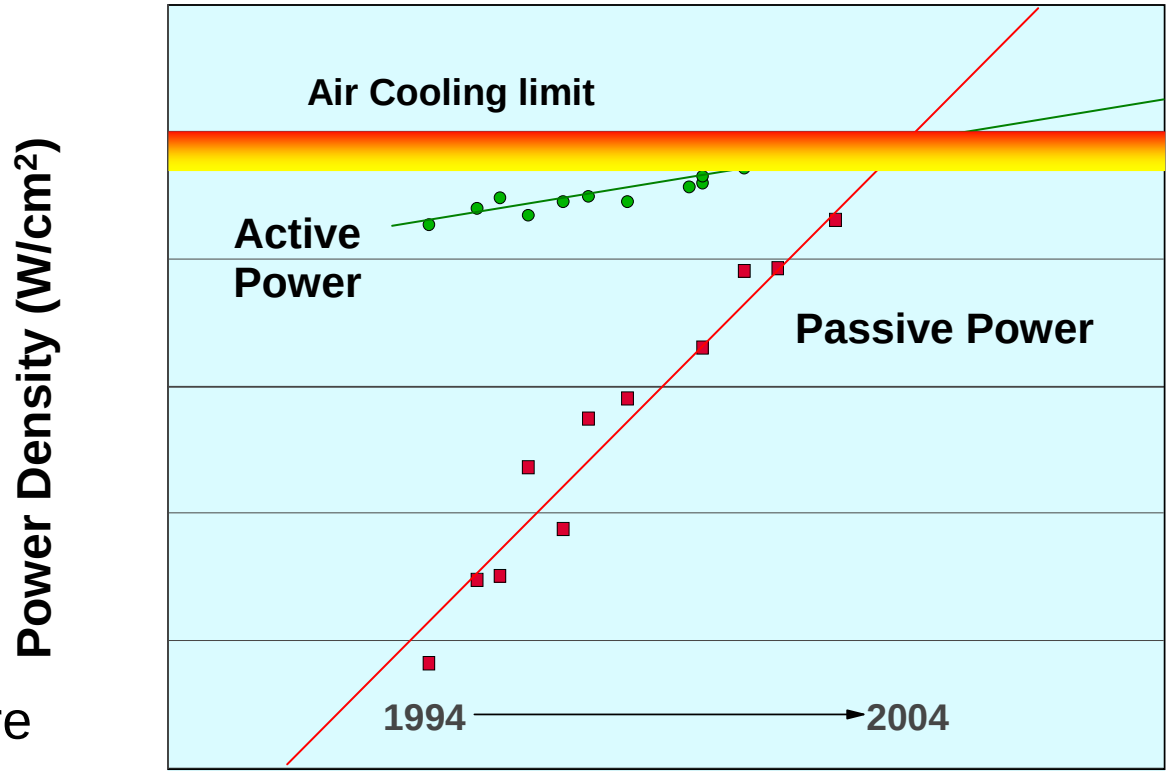

Better Performance Without Scaling

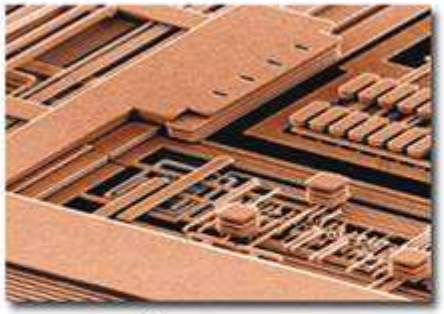

Copper

SOI
(Silicon-on-Insulator)

SiGe

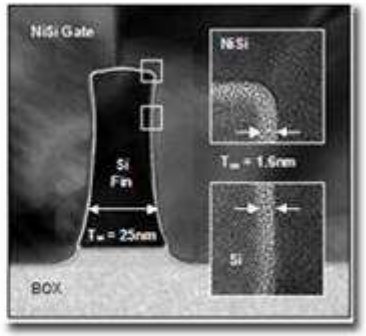

FinFET Double Gate

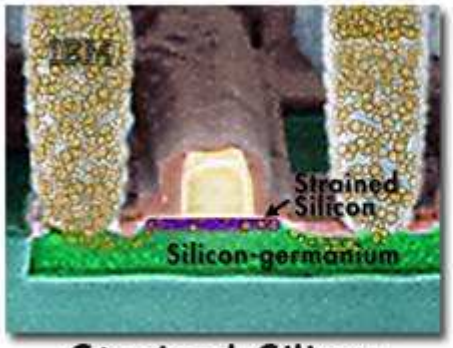

Strained Silicon

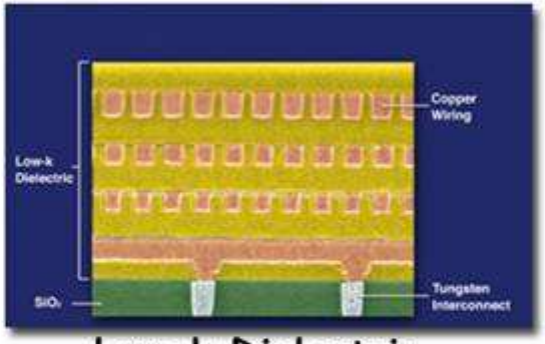

Low-k Dielectric

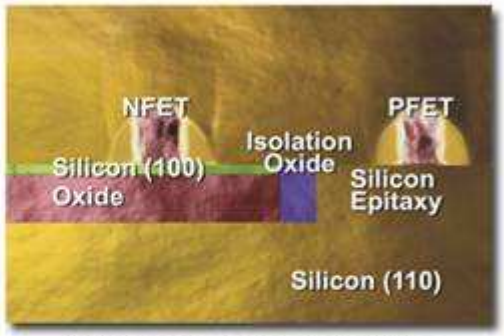

Hybrid-Orientation
Technology (HOT)

New materials & structures are critical for continuing CMOS technology density and performance path, while mitigating power dissipation

Computing Paradigm Shift

Today:

- **Single thread performance hitting limits**
  - Architecture and process technology saturated
  - Small percentage gains expected to remain

But:

- **Signs of paradigm shift to application specific system customization**
  - Large multiple gains for specific applications
  - Cell
    - 50x on TRE, 120x on FFT
  - Datapower
    - XML acceleration
  - Many examples in embedded markets

Future:

- **Greater performance demands**
  - Immersive Interaction
    - 3D, real-time, gaming inspired applications
    - Rich media, data-intensive content
  - Sensory Computing
    - New network tier
    - Autonomous agents performing intelligent analysis on streaming data
      - A&D: battlefield coordination

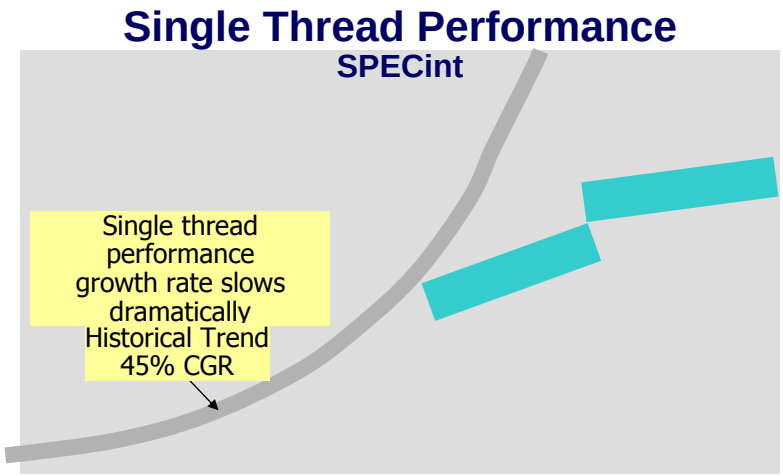

Solutions

- **Memory wall:**
  - More slower threads
  - Asynchronous loads
- **Efficiency wall:**
  - More slower threads
  - Specialized function
- **Power wall:**
  - Reduce transistor power
    - operating voltage
    - limit oxide thickness scaling
    - limit channel length
  - Reduce switching per function

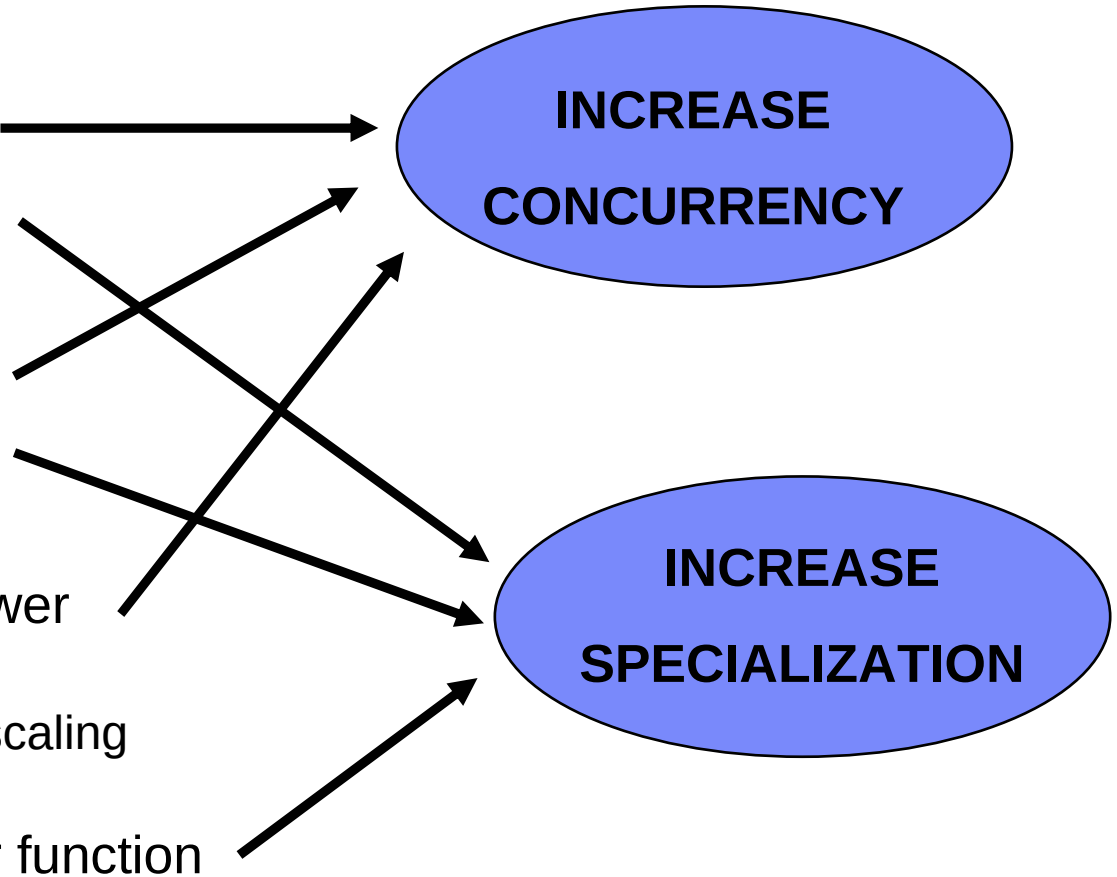

Architecture & Implementation

Cell Concept

- **Compatibility with 64b Power Architecture™**
  - Builds on and leverages IBM investment and community
- **Increased efficiency and performance**
  - Non Homogenous Coherent Chip Multiprocessor
    - Allows an attack on the “Frequency Wall”
  - Streaming DMA architecture attacks “Memory Wall”
  - High design frequency, low operating voltage attacks “Power Wall”
  - Highly optimized implementation
- **Interface between user and networked world**
  - Flexibility and security
  - Multi-OS support, including RTOS/non-RTOS
  - Architectural extensions for real-time management

Cell Architecture is …

64b Power Architecture™

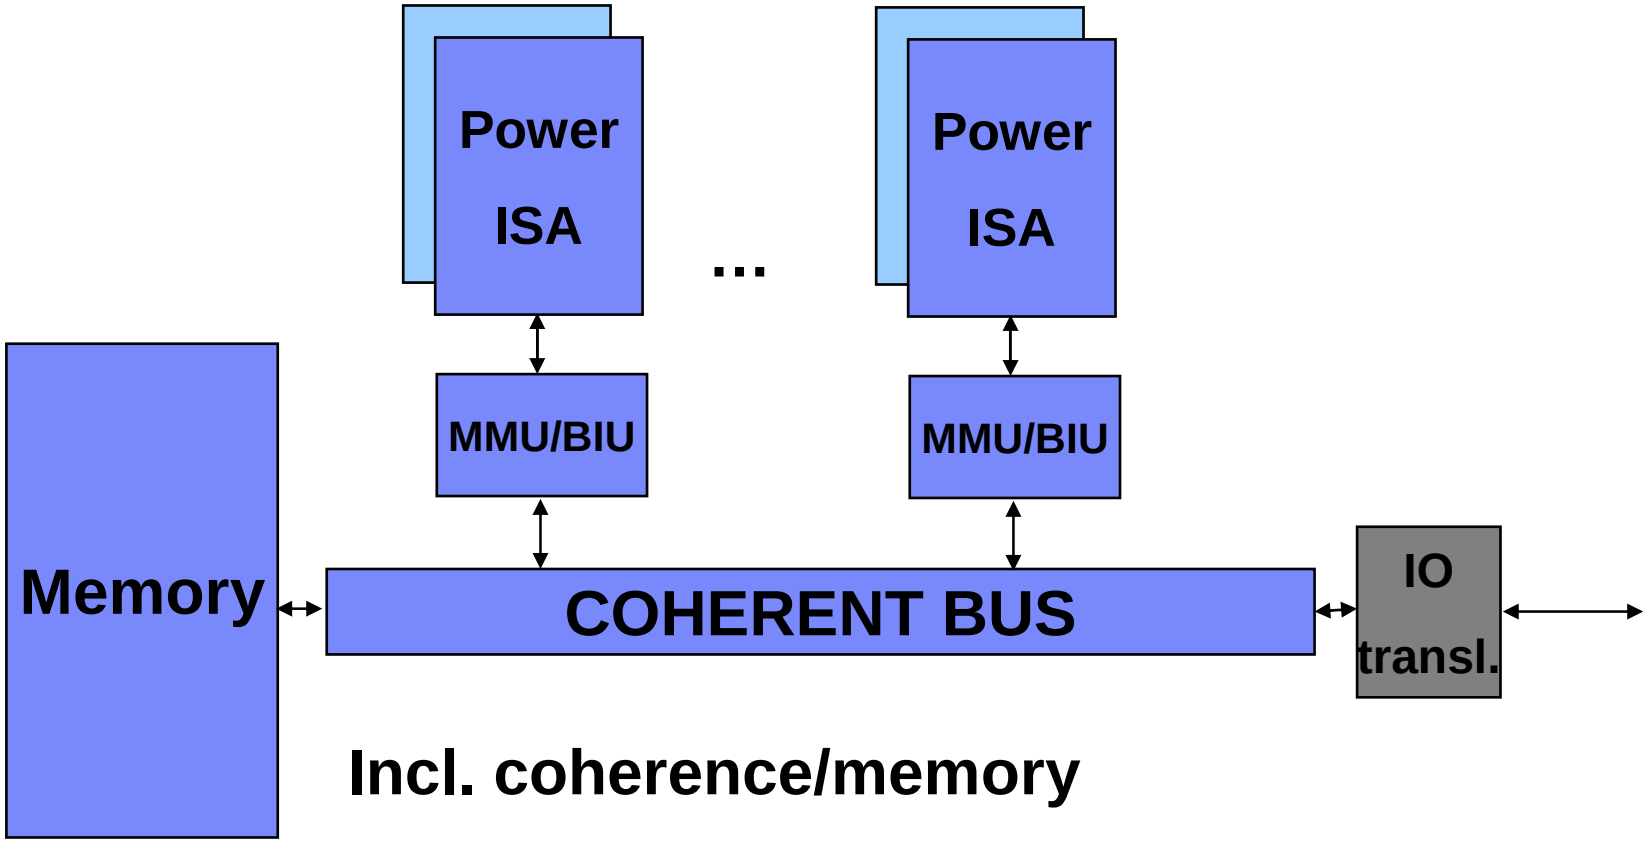

Incl. coherence/memory

compatible with 32/64b Power Arch. Applications and OS's

Cell Architecture is … 64b Power Architecture™

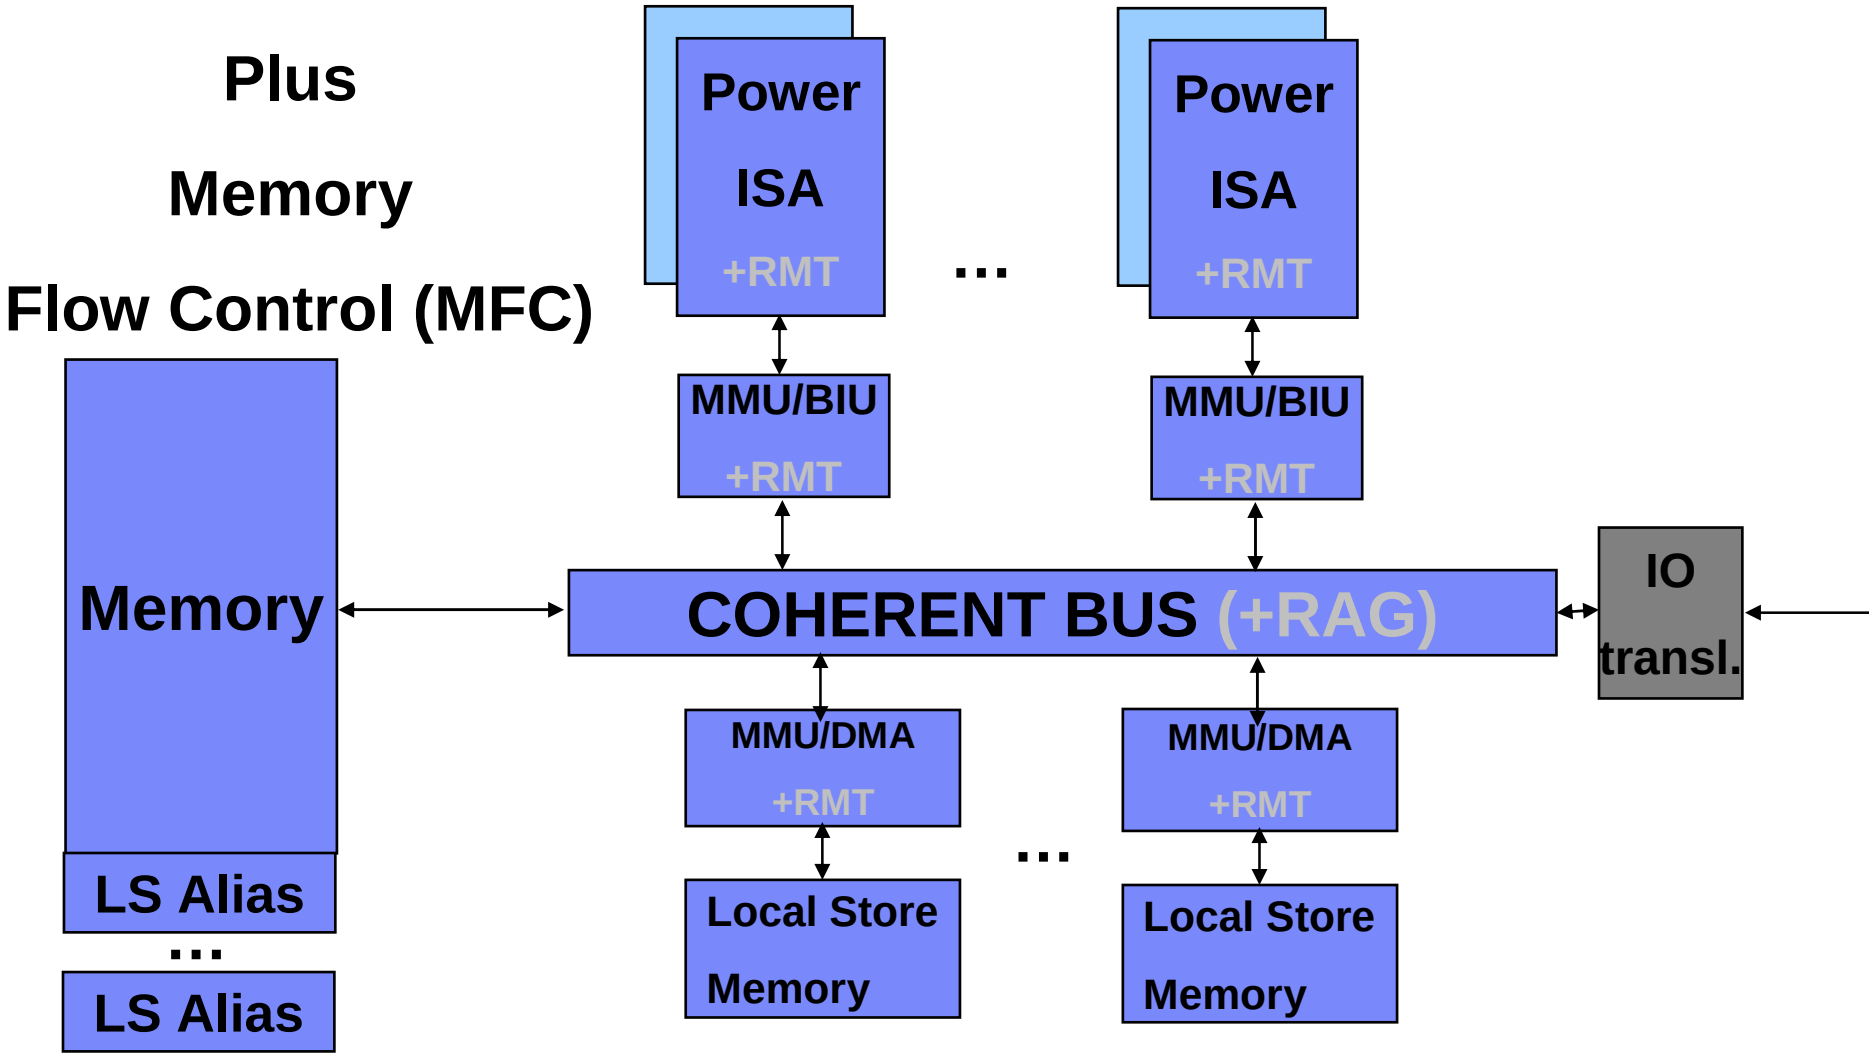

Cell Architecture is … 64b Power Architecture™+ MFC

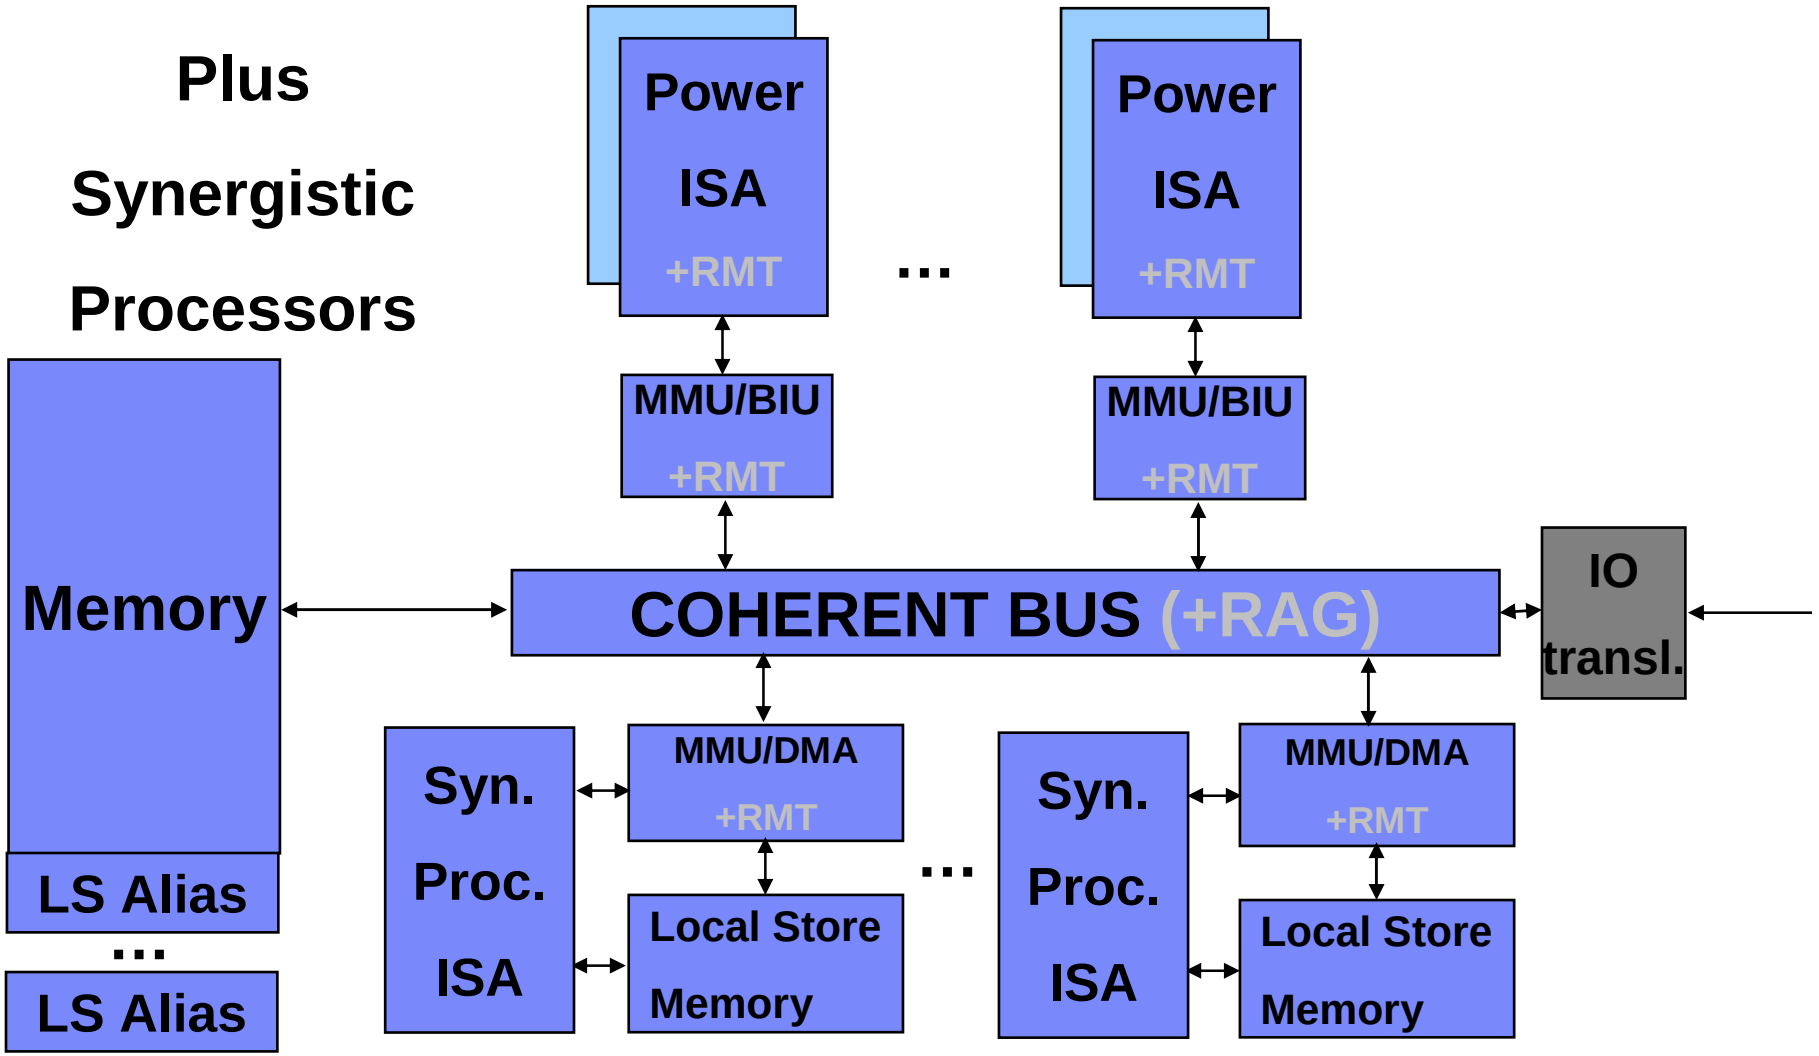

Coherent Offload Model

- **DMA into and out of Local Store equivalent to Power core loads & stores**
- **Governed by Power Architecture page and segment tables for translation and protection**
- **Shared memory model**
  - Power architecture compatible addressing
  - MMIO capabilities for SPEs
  - Local Store is mapped (alias) allowing LS to LS DMA transfers
  - DMA equivalents of locking loads & stores
  - OS management/virtualization of SPEs
    - Pre-emptive context switch is supported (but not efficient)

SPE Highlights

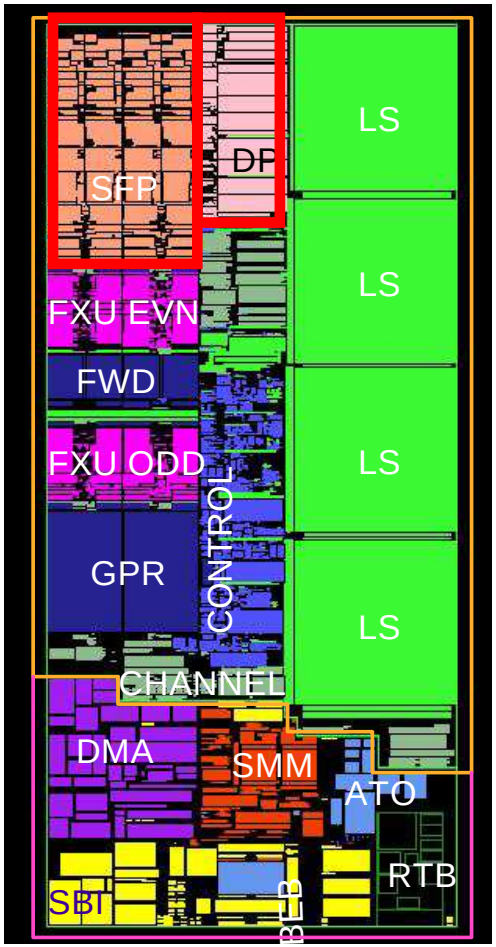

14.5mm² (90nm SOI)

- **RISC like organization**
  - 32 bit fixed instructions
  - Clean design – unified Register file
- **User-mode architecture**
  - No translation/protection within SPU
  - DMA is full Power Arch protect/x-late
- **VMX-like SIMD dataflow**
  - Broad set of operations (8 / 16 / 32 Byte)
  - Graphics SP-Float
  - IEEE DP-Float
- **Unified register file**
  - 128 entry x 128 bit
- **256KB Local Store**
  - Combined I & D

SPE BLOCK DIAGRAM

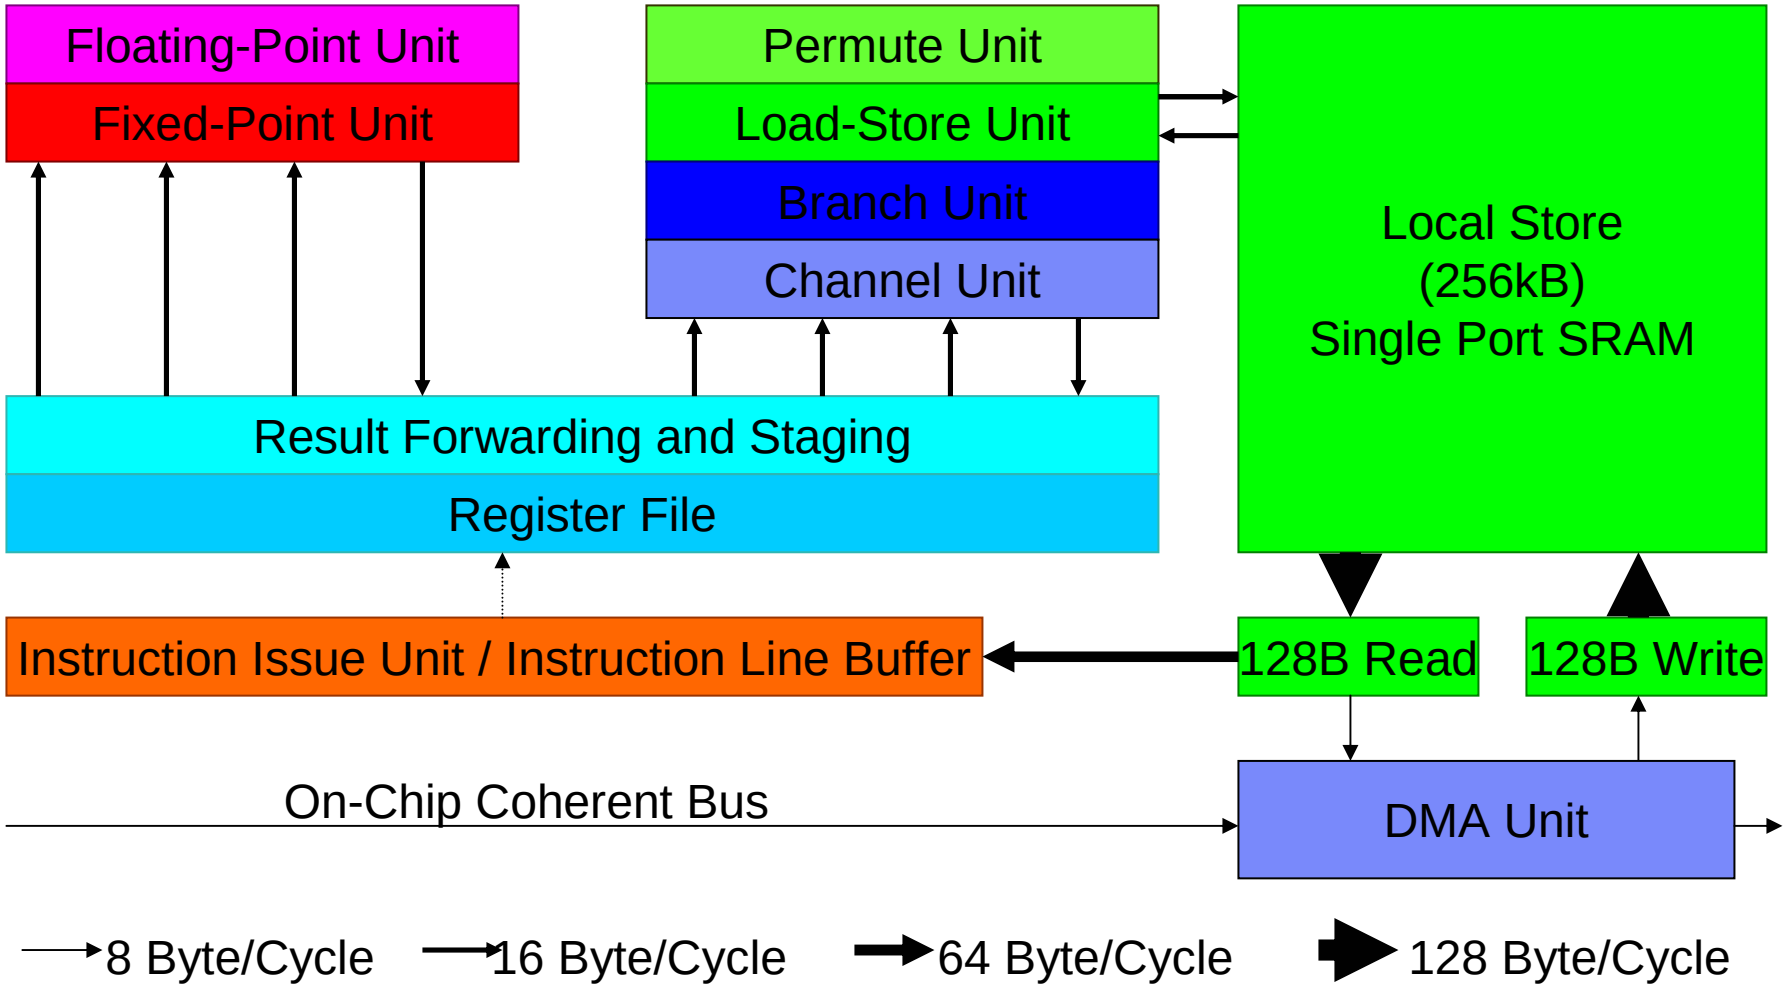

SPE PIPELINE FRONT END

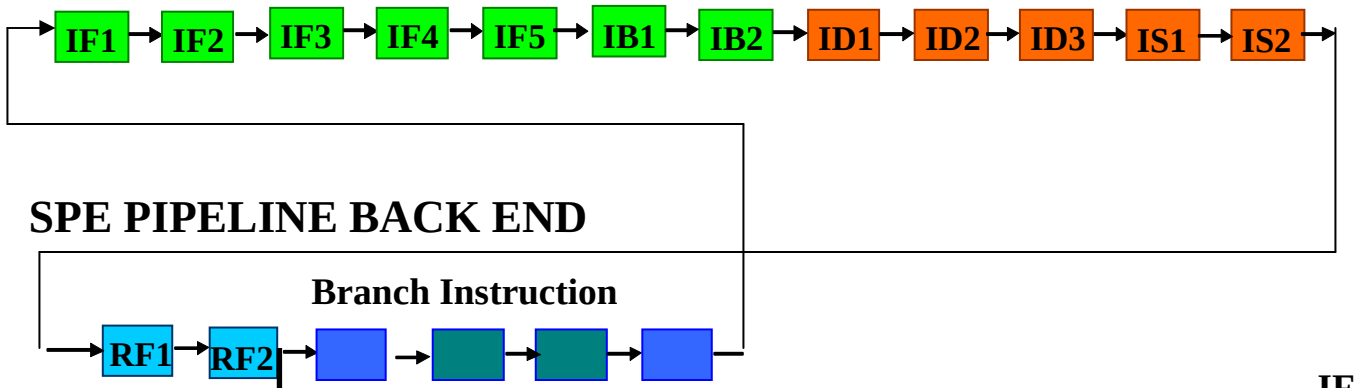

SPE PIPELINE BACK END

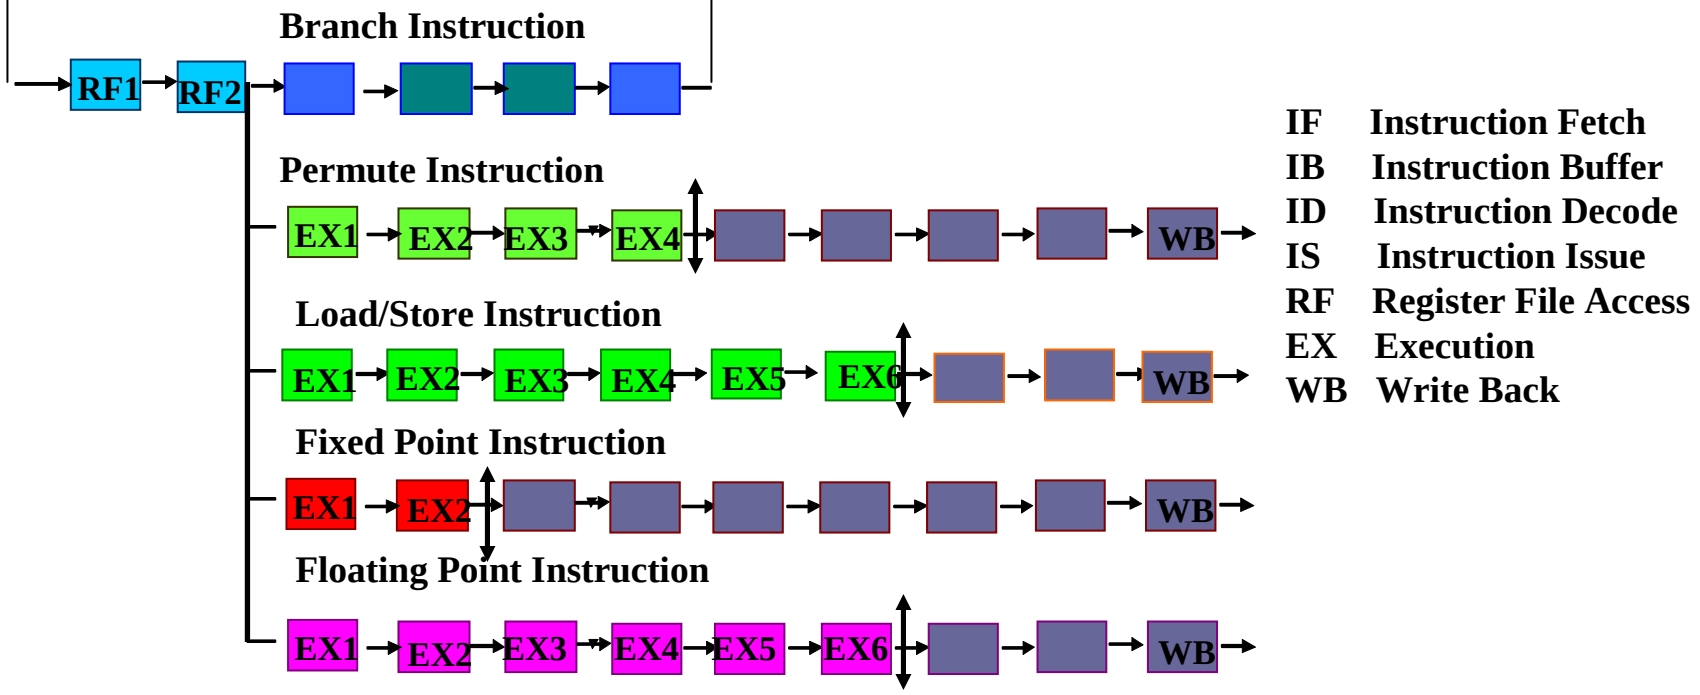

Cell Broadband Engine Processor

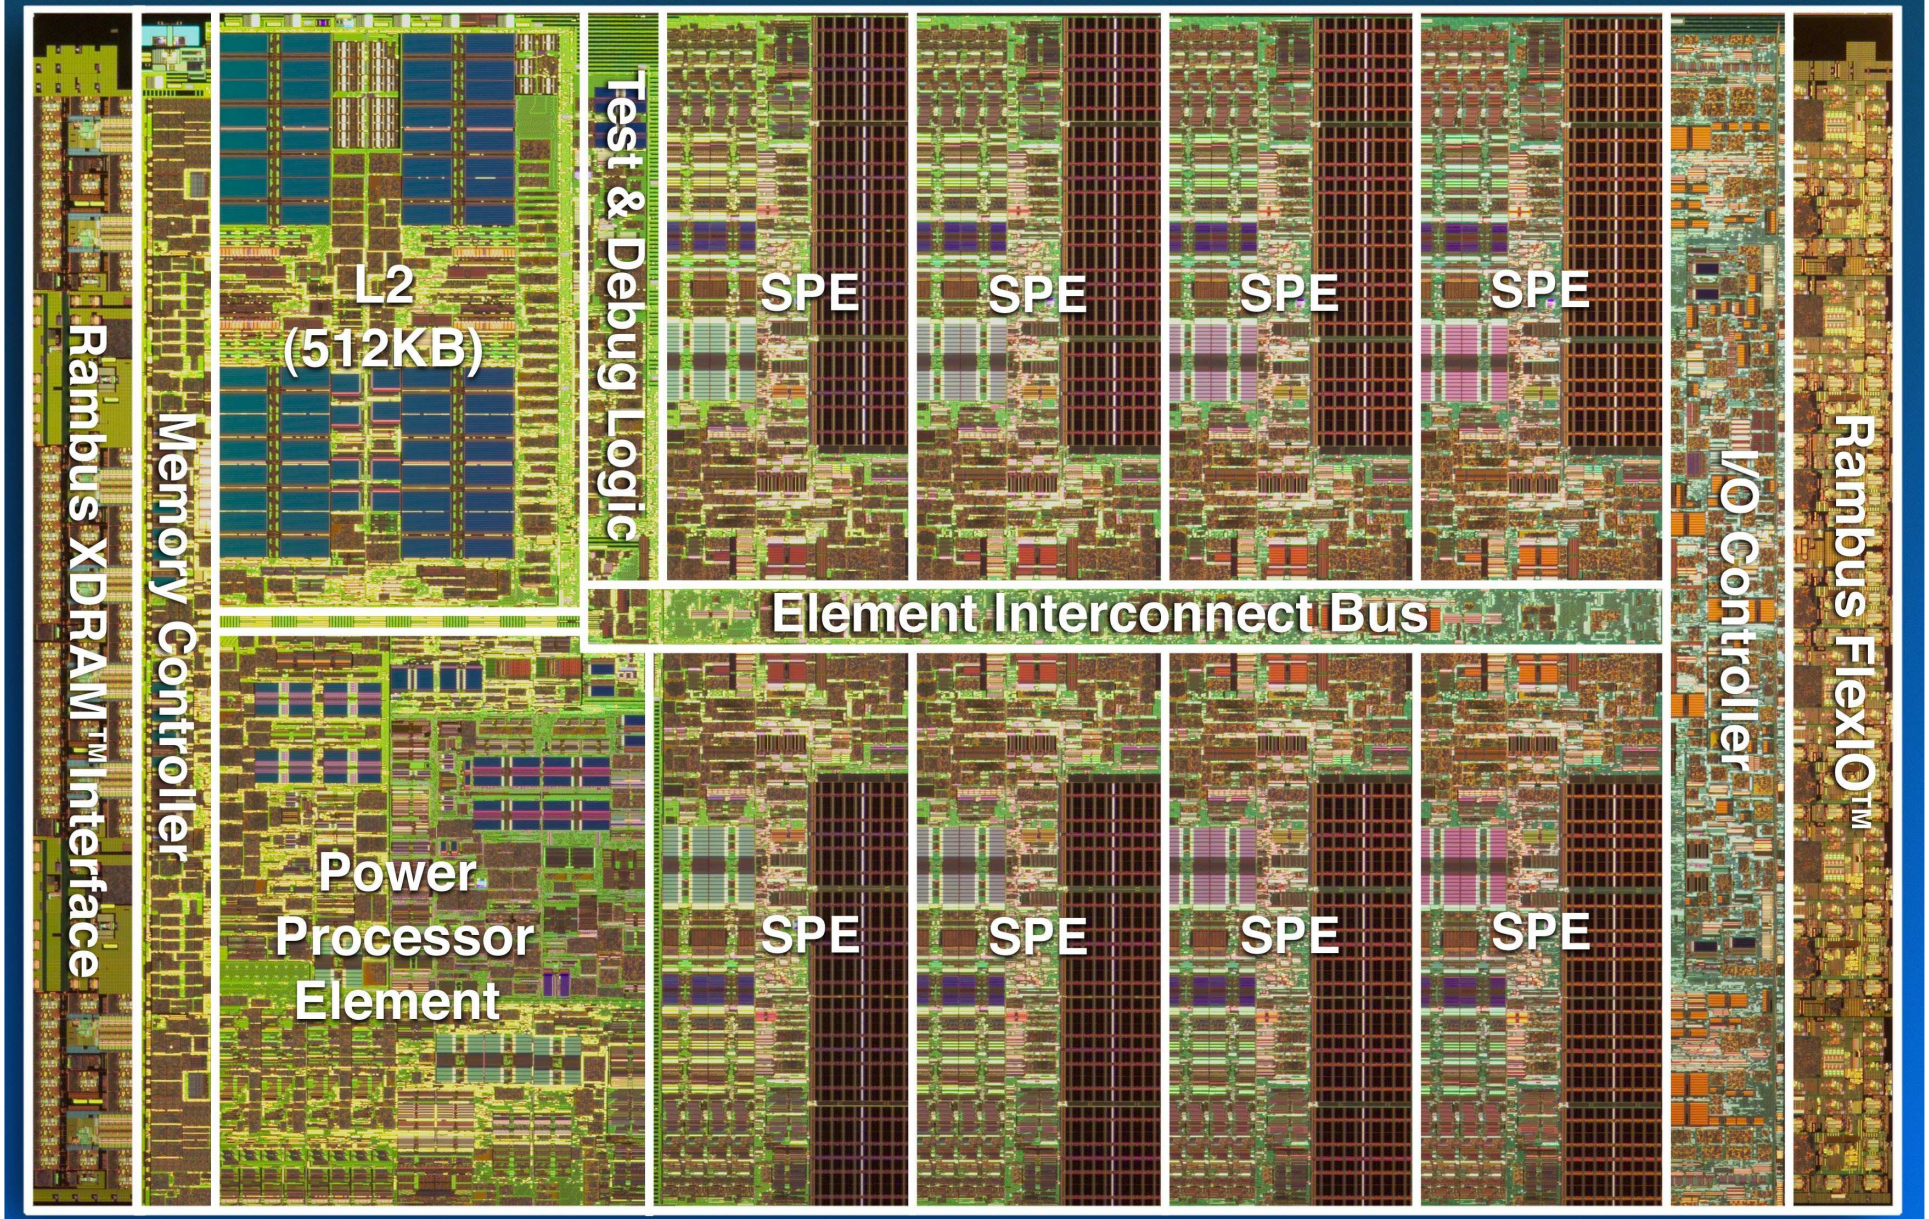

CELL PROCESSOR STATISTICS

- **250M transistors … 235mm2**
- **Top frequency >4GHz**
  - Lab conditions
  - Most efficient at ~1V
- **> 200 GFlops (SP) @3.2GHz**
- **> 20 GFlops (DP) @3.2GHz**
- **Up to 25.6 GB/s memory B/W**
- **Up to 70+ GB/s I/O B/W**
  - Practical ~ 50GB/s
- **100+ simultaneous bus transactions**
  - 16+8 entry DMA queue per SPE

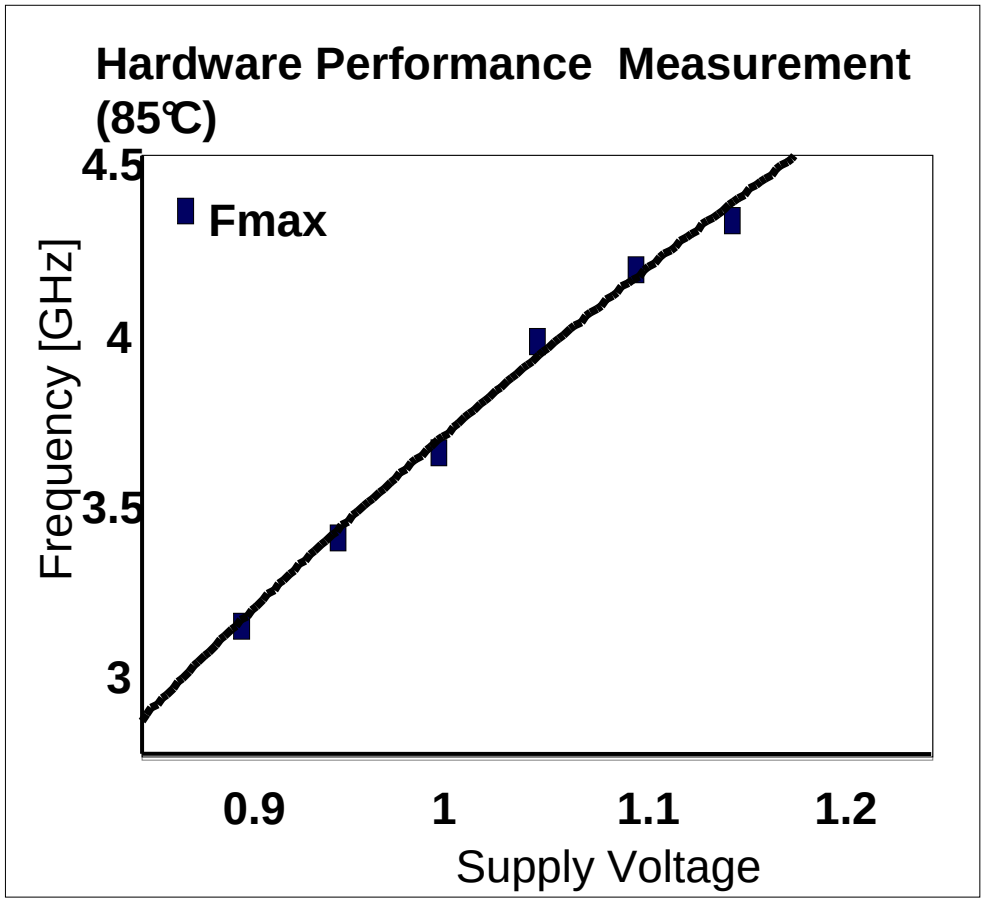

First pass hardware measurement in the Lab - Nominal Voltage = 1V

Programming

Cell Prototype Software Environment
www.ibm.com/developerworks/power/cell

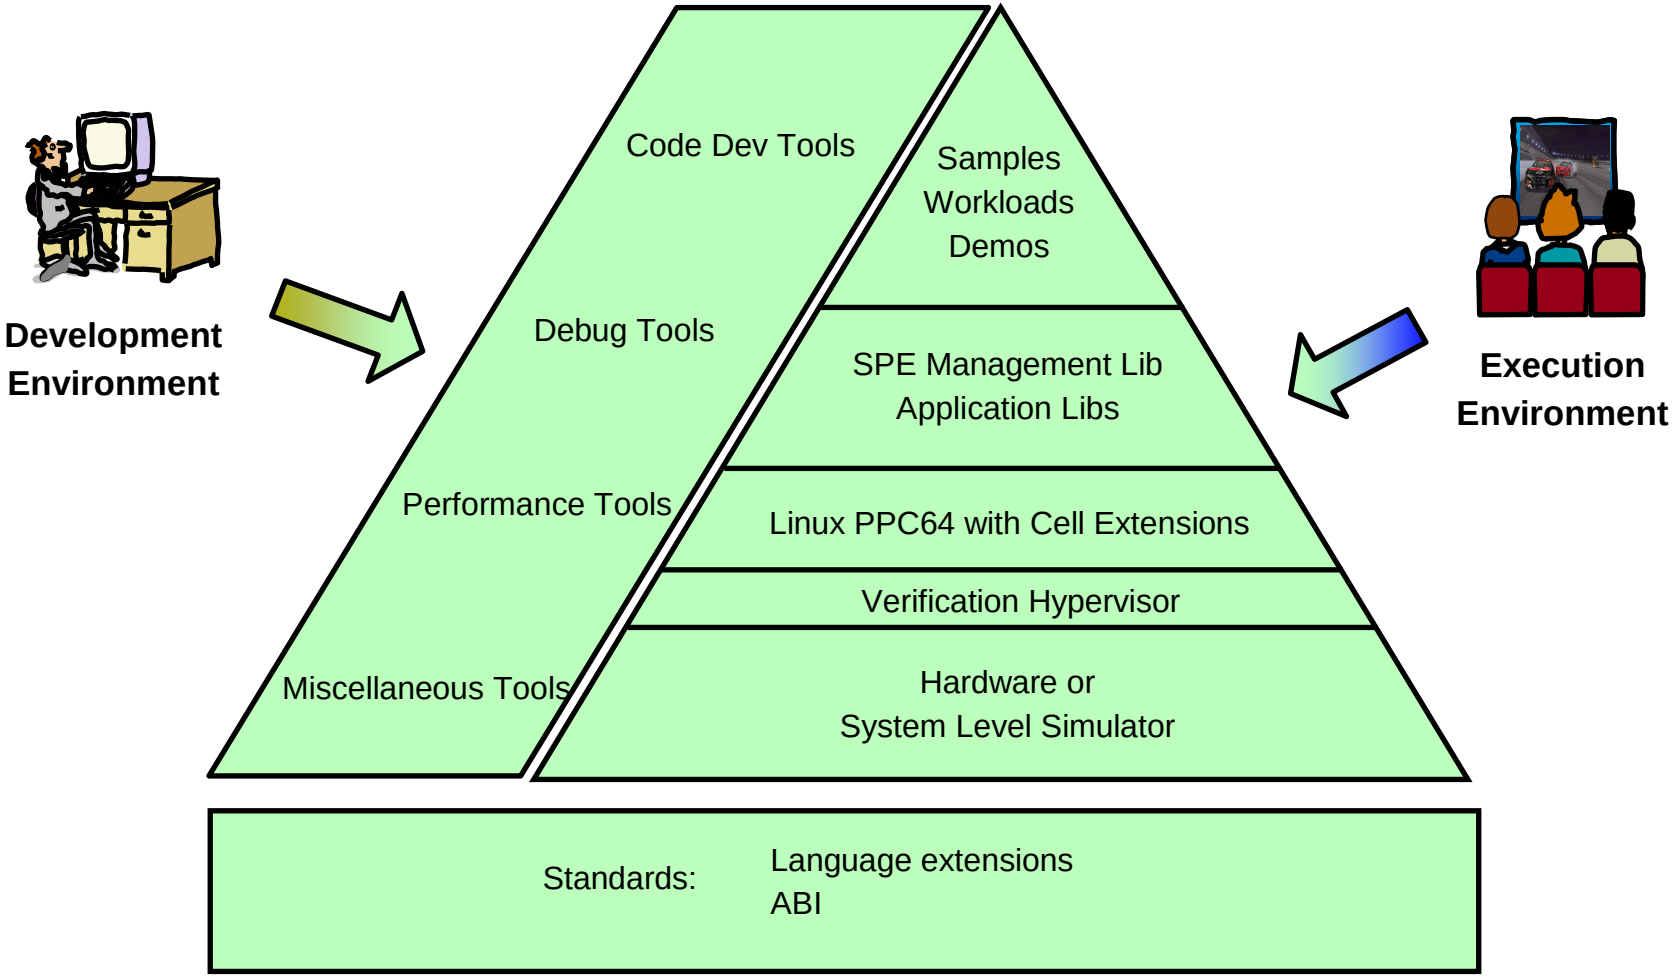

Operating System Runtime Strategy

- **Heterogeneous Multi-Threading Model**
  - PPE Threads, SPE Threads
  - SPE DMA EA = PPE Process EA Space
    - Or SPE Private EA space
  - OS supports Create/Destroy SPE tasks
  - Atomic Update Primitives used for Mutex
  - SPE Context Fully Managed
    - Context Save/Restore for Debug
    - Virtualization Mode (indirect access)
    - Direct Access Mode (realtime)
  - OS assignment of SPE threads to SPEs
    - Programmer directed using affinity mask
  - SPE Compilers use OS runtime services

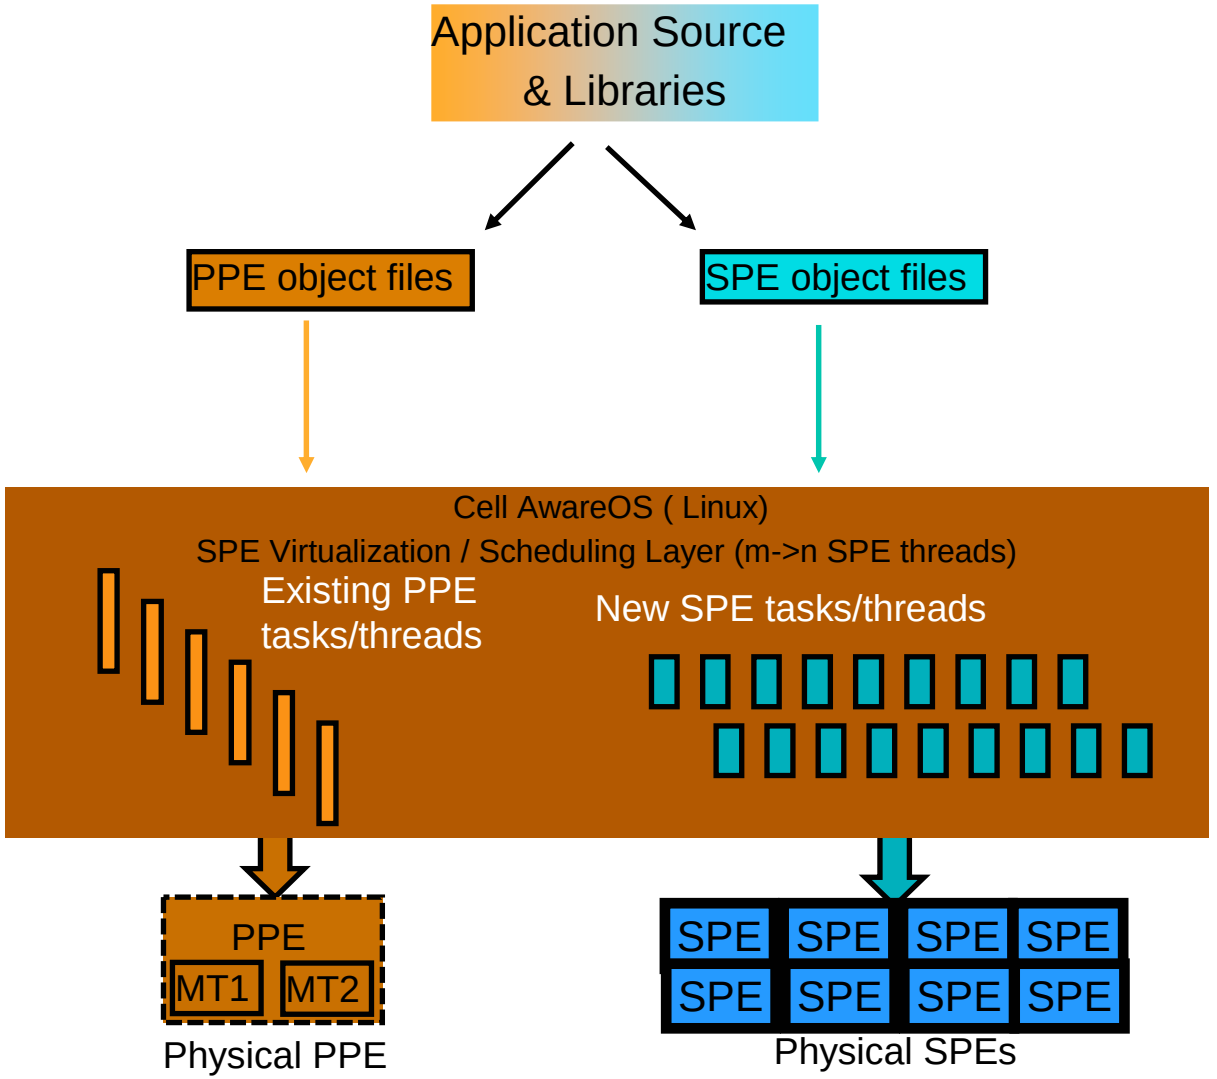

Programming Models

1) Application Specific Accelerators

Acceleration provided by O/S services

Application independent of accelerators platform fixed

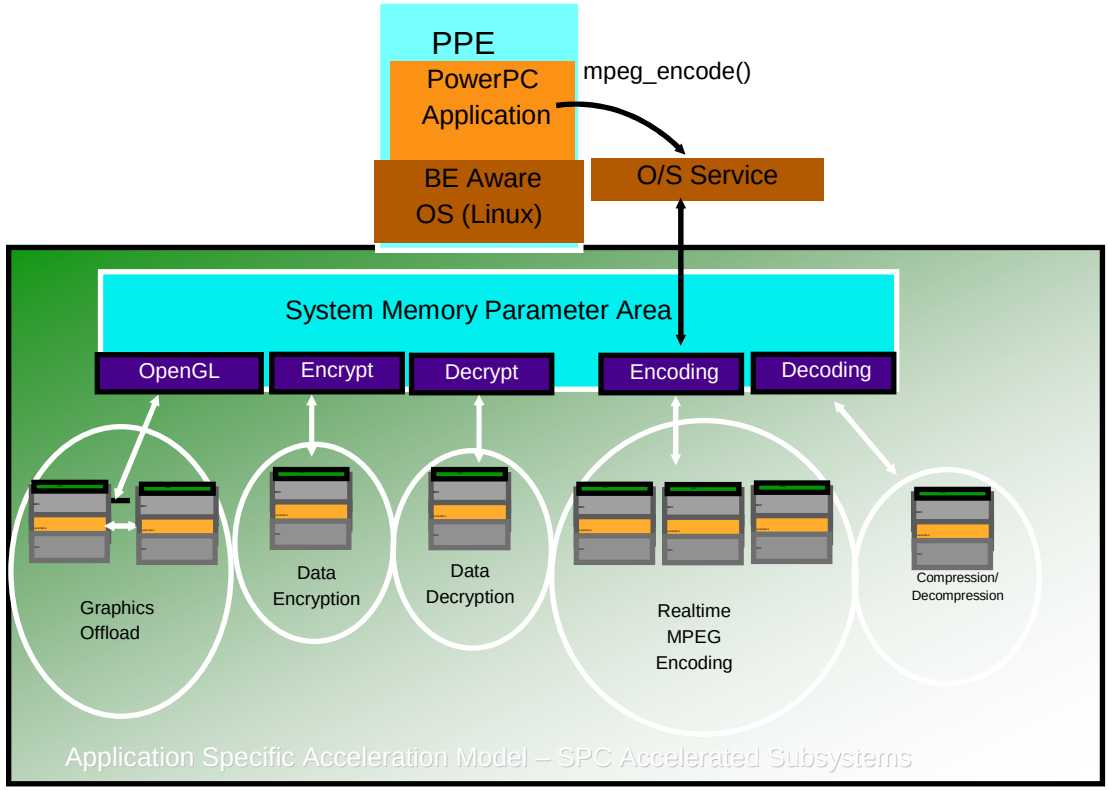

Programming Models

2) Function Offload

- SPE function provided by libraries
- Predetermined functions
- Application calls standard Libraries
  - Single source compilation
  - SPE working set fits in Local Store
  - O/S handles SPE allocation

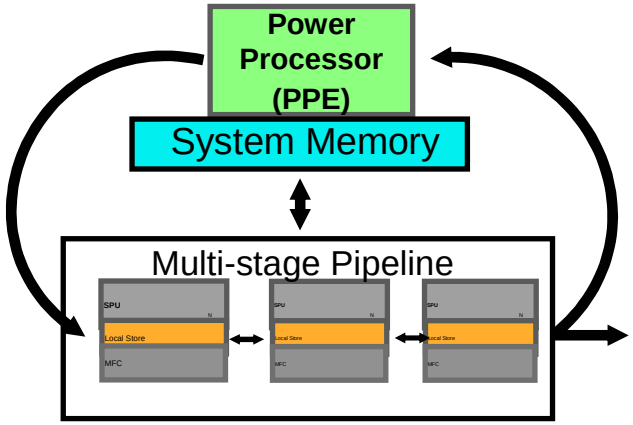

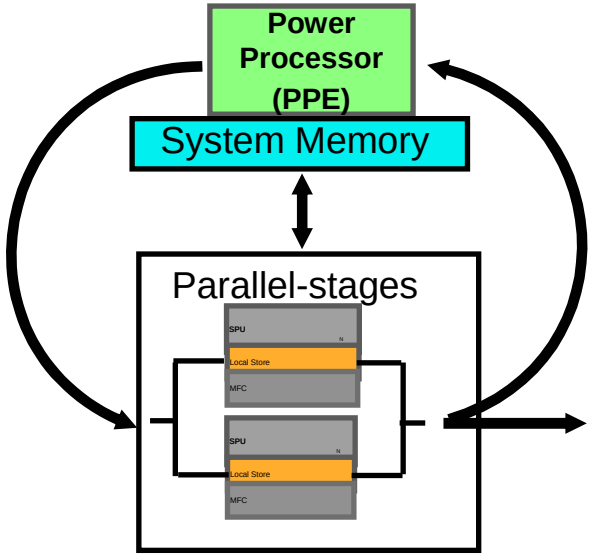

Programming Models

3) Computational Acceleration

- User created RPC libraries
  - User acceleration routines
  - User compiles SPE code
- Local Data
  - Data and Parameters passed in call
- Global Data
  - Data and Parameters passed in call
  - SPE Code manages global data

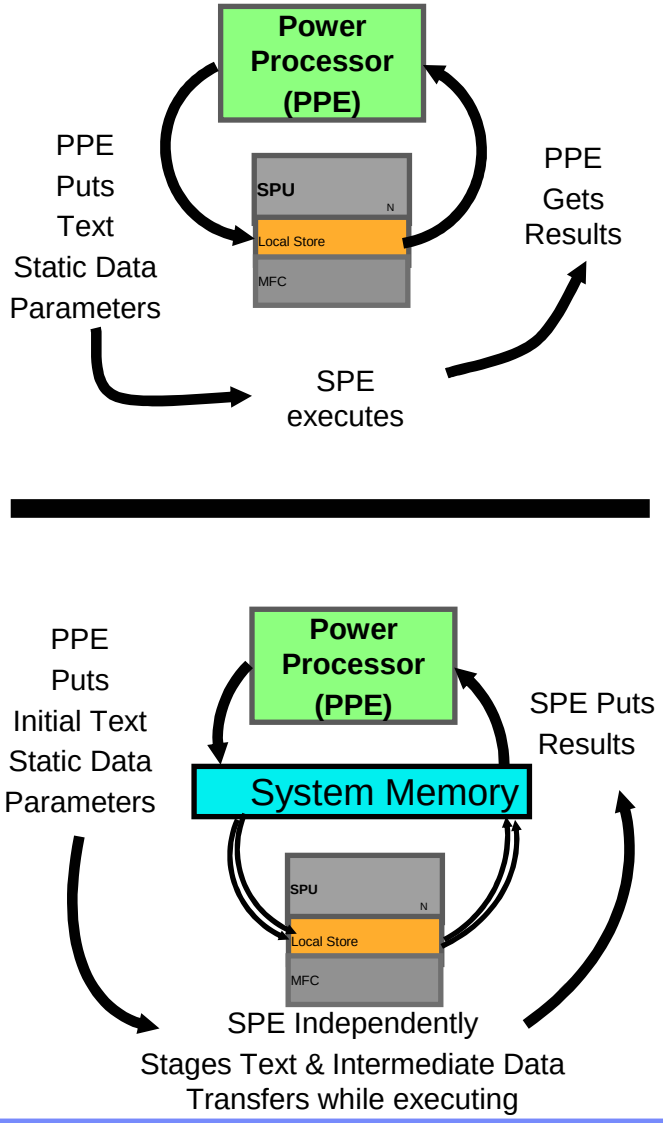

Single source approach to programming Cell

- **Single Source Compiler**
  - Auto parallelization (treat target Cell as an SMP)
  - Auto SIMD-ization (SIMD-vectorization)
  - Compiler management of Local Store as 2nd level register file / SW managed cache (I&D)
    - Most Cell unique piece
- **Optimization**
  - OpenMP pragmas
  - Vector.org SIMD intrinsics
  - Data/Code partitioning
  - Streaming / pre-specifying code/data use
- **Prototype Single Source Compiler Developed in IBM Research**

Applications

Cell BE Performance Characteristics

VERY GOOD

- **Computationally intensive code (loops can be unrolled)**
  - Order of magnitude more flops
- **Scatter-gather type problems**

  (e.g.FFT, Raycasting, sparse matrices(?))
  - Almost two orders of magnitude more performance than a typical PC processor

NOT OPTIMIZED FOR

- **Rapid context switch**
  - LS is context, switch is about 30uSec
  - Run to completion is preferred
  - Cooperative switching works well
  - Pre-emptive switch is possible
- **Load-compare-add-branch (TPCC, gcc type codes)**
  - 6 cycle load hurts
  - Still 8 (10) threads on a single chip

Terrain Rendering Engine

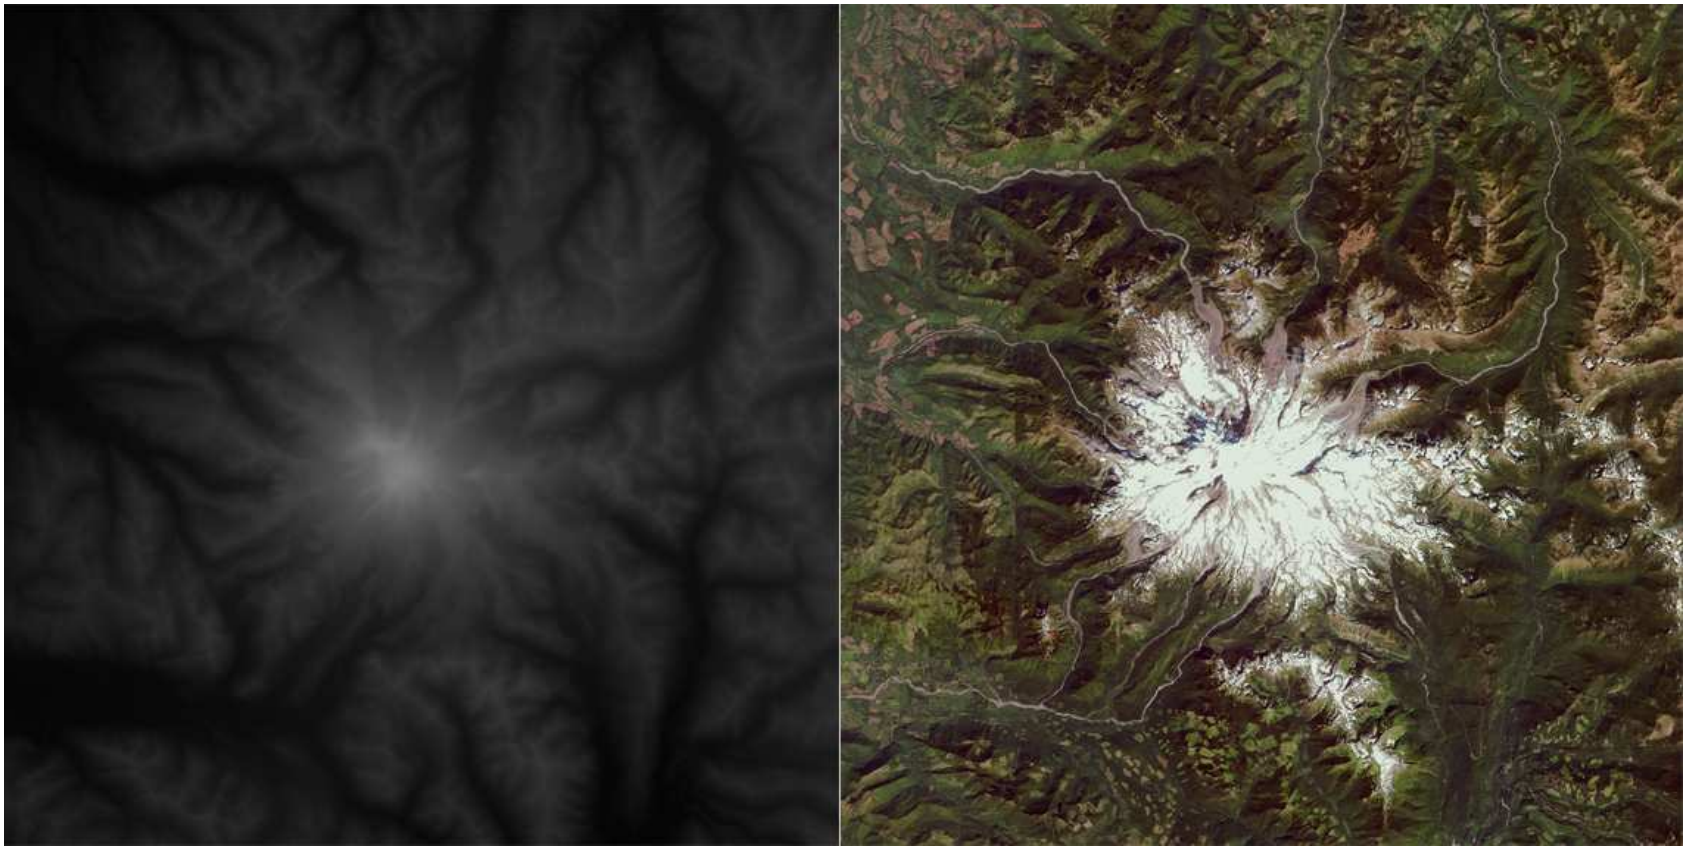

Planned Usage of Cell

- **Sony**
  - **Playstation 3**
    - 10 – 20 million units / year
    - Current Sony installed base 190 million units (PS1 & PS2)
  - **Cell Development Center reporting directly to CEO**

- **Toshiba**
  - **Cell on PC-type form factor card (October 2005)**
    - White box form factor
    - Reference system for Cell

- **IBM**
  - Engineering and Technology Services
    - Custom designs and application utilizing cell
    - Mercury Computer
  - Cell based systems

User Interaction Drives Innovation in Computing

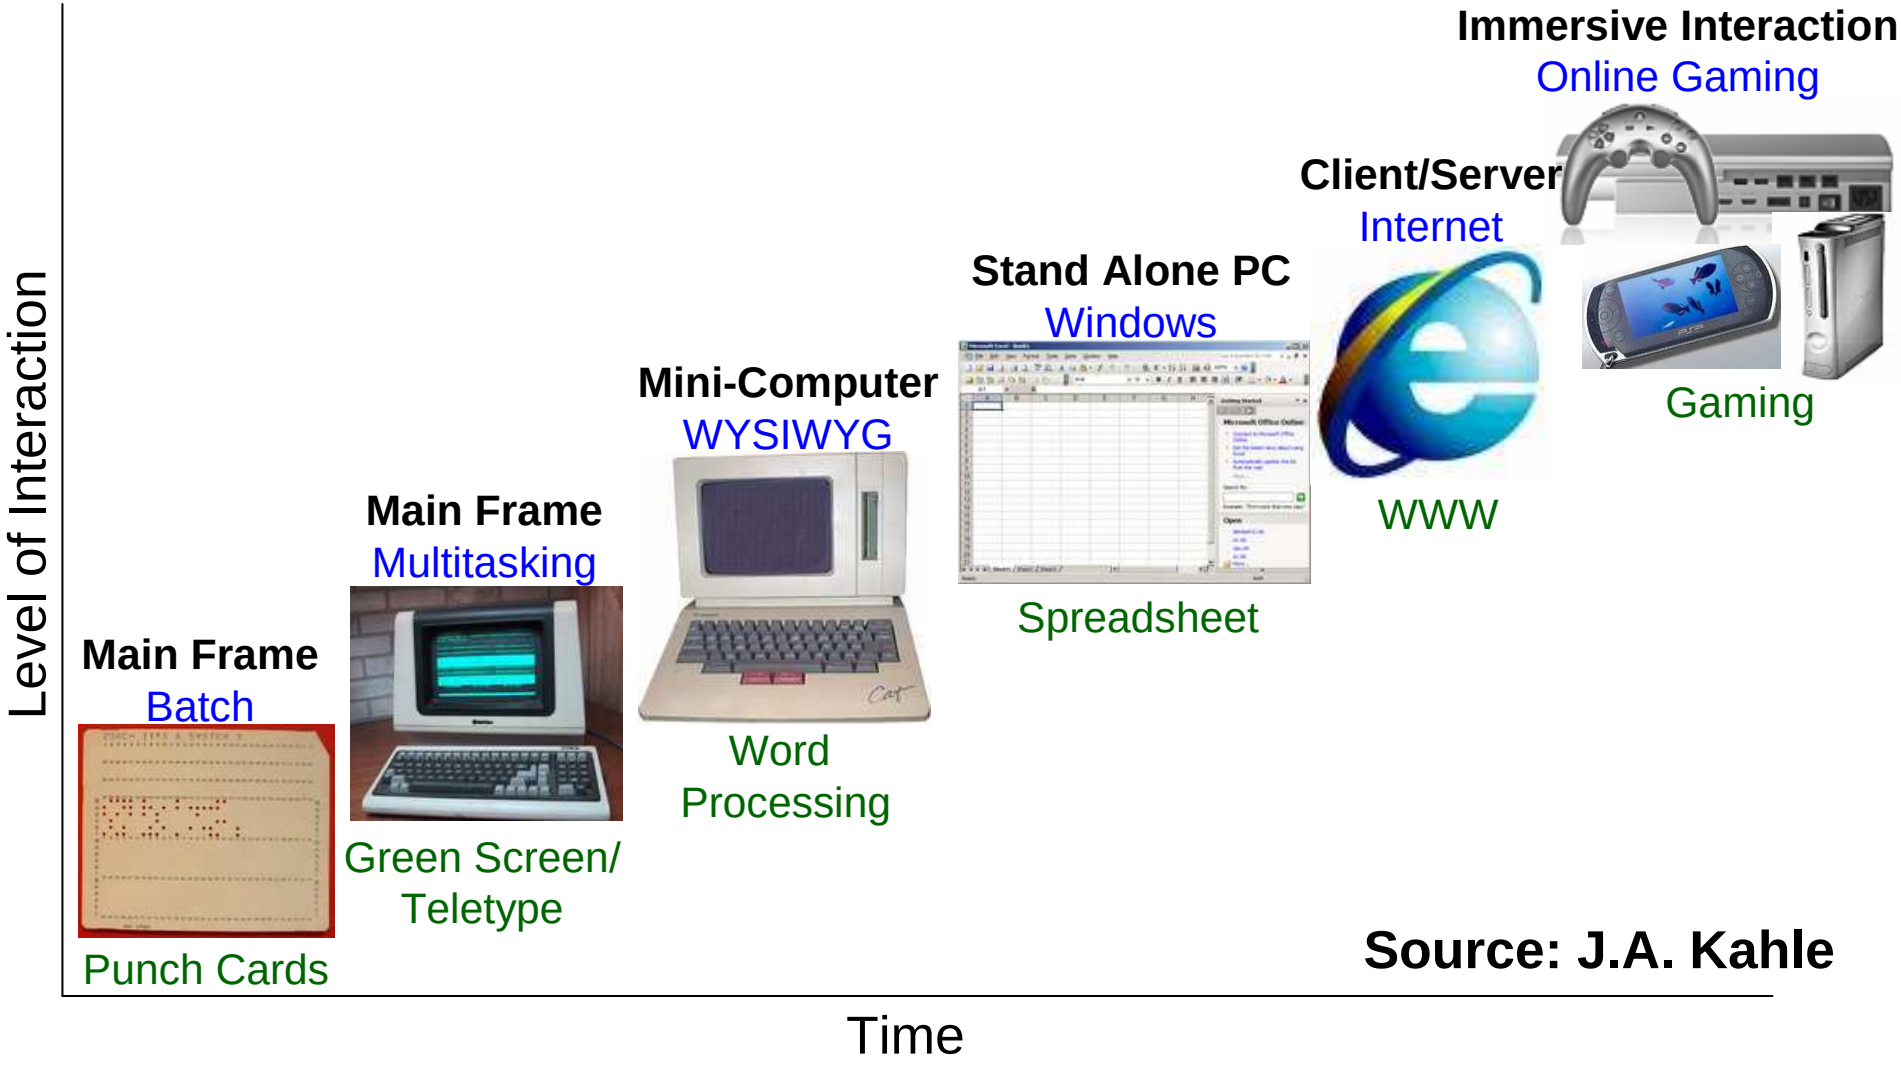

Characteristics of the Latest Transition in User Interaction

- Windows
- Click and wait…
- Client-centric
- User data accessible from client only
- Device-centric
- Connected
- Wired, sporadic
- E-mail/newsgroups

- Immersive, 3D interactivity
- Real-time
- Distributed
- User data accessible everywhere
- Device-agnostic
- Collaborative
- Wireless, always-on
- Text messaging/blogs

Summary & Conclusions

Summary

- **Cell ushers in a new era of leading edge processors optimized for digital media and entertainment**
- **Desire for realism is driving a convergence between supercomputing and entertainment**
- **New levels of performance and power efficiency beyond what is achieved by PC processors**
- **Responsiveness to the human user and the network are key drivers for Cell**
- **Cell will enable entirely new classes of applications, even beyond those we contemplate today**
- **TIME TO GET IN THE GAME!**